UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS

NUANCE COMMUNICATIONS, INC., Plaintiff, v. OMILIA NATURAL LANGUAGE SOLUTIONS, LTD., Defendant.	Civil Action No. 19-11438-PBS

MEMORANDUM AND ORDER ON CLAIM CONSTRUCTION

August 6, 2020

Saris, D.J.

Nuance Communications, Inc. accuses Omilia Natural Language Solutions, Ltd., of infringing U.S. Patent No. 6,999,925 (the "'925 patent"). The parties dispute the claim construction of four terms: "a second language," "multi-lingual speech recognizer," "generating a second acoustic model," and "automatically generate/ing." The Court held a non-evidentiary Markman hearing on July 10, 2020 and reviewed technical tutorials and briefs submitted by both parties.

BACKGROUND

The '925 patent, entitled "Method and Apparatus for Phonetic Context Adaptation for Improved Speech Recognition," describes a method to "automatically generat[e] from a first speech recognizer a second speech recognizer which can be

adapted to a specific domain." '925 patent, col. 1, ll. 15-19 (Dkt. 84-1). The patented method "provide[s] for fast and easy customization of speech recognizers to a given domain," such as "a certain language, a multitude of languages, a dialect or a set of dialects, [or] a certain task area" like medicine or banking. See id. at col. 2, ll. 21-23, col. 6, ll. 5-8. Speech recognizers generally match audio input of human speech to phonemes, or abstractions of individual sounds. To do so, recognizers determine the probability that an audio file contains a certain sound, or phone, using the phone's "phonetic context." "Phonetic context" refers to sub-word units consisting of individual phones with neighboring phones. See 'id. at col. 4, ll. 35-38.

"[T]he collection of a large amount of training data and the subsequent training of a speech recognizer is both expensive and time consuming." Id. at col. 1, ll. 61-63. A recognizer is trained using "labelled training data." See id. at col. 1, ll. 54-55. The recognizer uses the training data to "construct[] a binary decision network" with a "predefined number of leaves," which are also known as "terminal nodes." See id. at col. 4, ll. 44, 61, col. 5, l. 6.

Several methods for "the customization of a general speech recognizer to a particular domain" predate the '925 patent. See id. at col. 5, ll. 28-67. These methods included "[1]

'selecting' a subset of the general speech recognizer decision network and the phonetic contexts or [2] simply 'enhancing' the decision network by . . . attaching a new sub-tree with new leaf nodes and further phonetic contexts." Id. at col. 7, ll. 12-17. The '925 patent describes the first approach as sub-optimal because it could not "detect any new phonetic context . . . relevant to a new domain but not present in the general recognizer's inventory." Id. at col. 5, ll. 59-61. The patent describes the second approach as sub-optimal because it "still require[d] the collection of a substantial amount" of training data. Id. at col. 5, ll. 50-51; see also id. at col. 1, l. 66-col. 2, l. 3 (describing a third "[c]onventional adaptation method[]" of "simply provid[ing] a modification of the acoustic model parameters")..

In contrast, the '925 patent method requires only a "small amount of domain specific adaptation data" to adapt a speech recognizer to a particular domain. Id. at col. 6, ll. 11-16. The patent "preserves the phonetic context information of the first speech recognizer" and "simultaneously allows for the creation of new phonetic contexts." Id. at col. 2, ll. 45-50. "This is achieved by . . . re-estimating the decision network and phonetic contexts based on domain-specific training data." Id. at col. 6, ll. 16-20.

Figure 1 reflects a preferred embodiment of the '925 patent. The specification makes clear "that the invention is not to limited to the precise arrangements and instrumentalities shown" in Figure 1. Id. at col. 2, ll. 63-64.

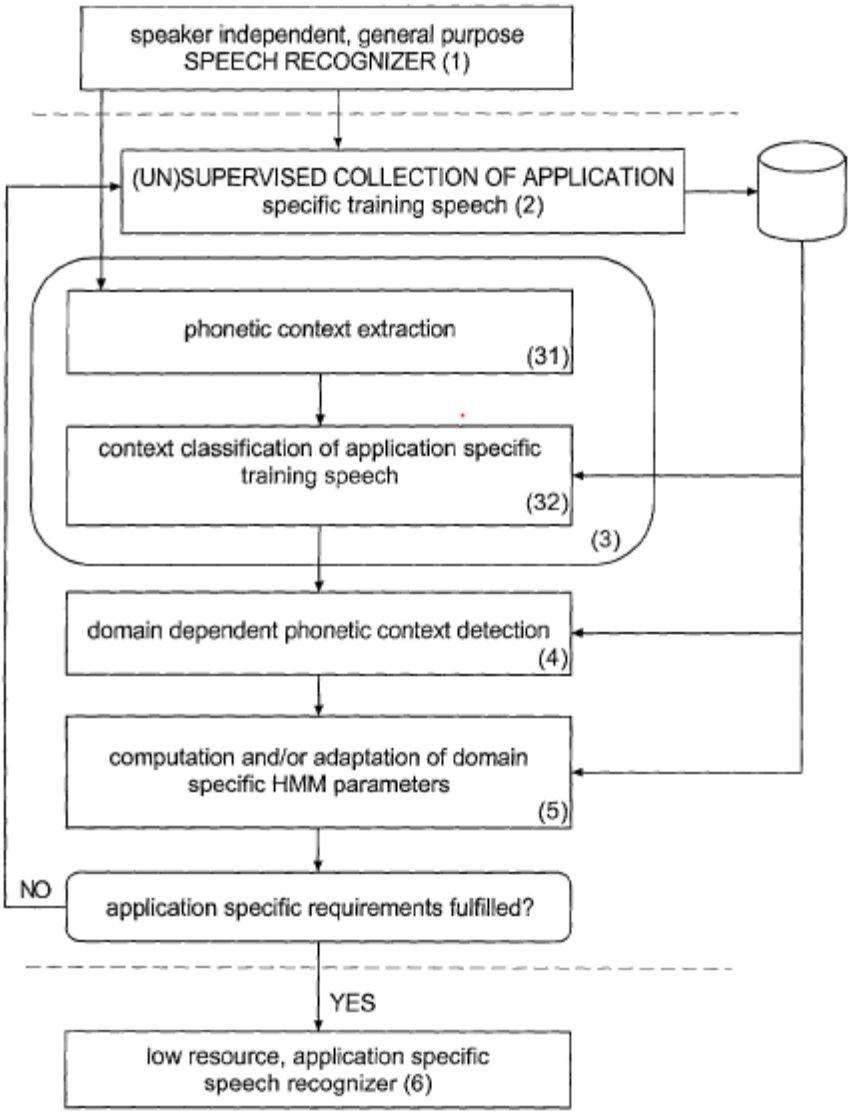

FIGURE 1

The patent describes the embodiment in Figure 1 as follows. First, the adaptation data is "passed through the original decision network" to "obtain a partitioning of the adaptation data." Id. at col. 7, ll. 52-59. The system then "insert[s], delet[es], or adapt[s]" phonetic contexts, "resulting in a new, re-estimated (domain specific) decision network." Id. at col. 7, ll. 8-12, 62-65.

This method can "significantly improve the recognition rate within a given target domain" while "avoid[ing] an unacceptable decrease of recognition accuracy in the original recognizer's domain." Id. at col. 10, ll. 13-18. The patent provides that this method "can be used for the incremental and data driven incorporation of a new language into a true multi-lingual speech recognizer." Id. at col. 9, ll. 20-22.

Independent claims 1, 2, 14, 15, and 27 and dependent claims 12 and 25 of the '925 patent use the disputed terms at issue here. The claims are as follows:

> 1. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said method comprising:
>
> based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second

decision network utilize a phonetic decision [t]ree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, and wherein said re-estimating comprises partitioning said training data using said first decision network of said first speech recognizer.

2. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model wit[h] a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said method comprising:

based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, wherein said domain-specific training data is of a limited amount, and wherein the generating step further comprises the steps of:

identifying at least one acoustic context from the domain-specific training data; and

adding a node to the second decision network for the identified context independent of other generating step operations.

12. The method of claim 6, wherein said first speech recognizer is a speech recognizer of at least a first language and said domain specific training data relates to a second language and said second speech recognizer is a multi-lingual speech recognizer of said second language and said at least first language.

14. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision

network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said machine-readable storage causing the machine to perform the steps of:

based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding said phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, and wherein said re-estimating comprises portioning said training data using said first decision network of said first speech recognizer.

15. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said machine-readable storage causing the machine to perform the steps of:

based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding said phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, wherein said domain-specific training data is of a limited amount, and wherein the generating step further comprises the steps of:

identifying at least one acoustic context from the domain-specific training data; and

adding a node to the second decision network for the identified context independent of other generating step operations.

25. The machine-readable storage of claim 19, wherein said first speech recognizer is a speech recognizer of at least a first language and said domain specific training data relates to a second language and said second speech recognizer is a multi-lingual speech recognizer of said second language and said at least first language.

27. A computerized method of generating a second speech recognizer comprising the steps of:

identifying a first speech recognizer of a first domain comprising a first acoustic model with a first decision network and corresponding first phonetic contexts;

receiving domain-specific training data of a second domain; and

based on the first speech recognizer and the domain-specific training data, generating a second acoustic model of said first domain and said second domain comprising a second acoustic model with a second decision network and corresponding second phonetic contexts, wherein the first domain comprises at least a first language, wherein the second domain comprises at least a second language, and wherein the second speech recognizer is a multi-lingual speech recognizer.

Id. at claims 1, 2, 12, 14, 15, 25, 27 (emphases added).

LEGAL STANDARD

Claim construction is an issue of law for the court. Markman v. Westview Instruments, Inc., 517 U.S. 370, 372 (1996). "The words of a claim are generally given their ordinary and customary meaning as understood by a person of ordinary skill in the art when read in the context of the specification and prosecution history." Nevro Corp. v. Bos. Sci. Corp., 955 F.3d 35, 43 (Fed. Cir. 2020) (citation omitted). "There are only two

exceptions to this general rule: 1) when a patentee sets out a definition and acts as his own lexicographer, or 2) when the patentee disavows the full scope of the claim term either in the specification or during prosecution." Golden Bridge Tech., Inc. v. Apple Inc., 758 F.3d 1362, 1365 (Fed. Cir. 2014) (citation omitted).

A person of ordinary skill in the art looks to "the words of the claims themselves, the remainder of the specification, the prosecution history, and extrinsic evidence concerning relevant scientific principles, the meaning of technical terms, and the state of the art." Phillips v. AWH Corp., 415 F.3d 1303, 1314 (Fed. Cir. 2005) (en banc) (quoting Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc., 381 F.3d 1111, 1116 (Fed. Cir. 2004)).

Courts begin claim construction with "the language of the claims themselves." Trs. of Columbia Univ. in City of N.Y. v. Symantec Corp., 811 F.3d 1359, 1362 (Fed. Cir. 2016). "[C]laims must be read in view of the specification, of which they are a part." Id. (citation omitted). The specification "is the single best guide to the meaning of a disputed term" and is usually "dispositive." Cont. Circuits LLC v. Intel Corp., 915 F.3d 788, 796 (Fed. Cir. 2019) (citation omitted). However, courts should be careful not to limit the claims to the specific embodiments in the specification. See Phillips, 415 F.3d at 1320.

The prosecution history "often inform[s] the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution." Id. at 1317. Examining prosecution history "ensure[s] that claims are not argued one way in order to maintain their patentability and in a different way against accused infringers." Aylus Networks, Inc. v. Apple Inc., 856 F.3d 1353, 1360 (Fed. Cir. 2017).

In general, extrinsic evidence is "less reliable than the patent and its prosecution history in determining how to read claim terms." Phillips, 415 F.3d at 1318.

DISCUSSION

I. "a second language" (Claims 12, 25, 27)

Claim Term	Nuance's Proposed Construction	Omilia's Proposed Construction
"a second language"	a language other than the first language of the first domain/speech recognizer	a second language not incorporated in the first acoustic model / speech recognizer

Nuance argues that "a second language" is any language other than the first language of the first domain or speech recognizer. In Nuance's view, the first domain may therefore include the first and second language. Omilia disagrees, arguing that the second language cannot be at all "incorporated" in the decision network of the first speech recognizer.

The Court begins its analysis with the claim language. Dependent claims 12 and 25 provide that the "first speech recognizer is a speech recognizer of at least a first language," "domain specific training data relates to a second language," and the "second speech recognizer is a multi-lingual speech recognizer of said second language and said at least first language." '925 patent at col. 11, ll. 59-64, col. 13, l. 29-col. 14, l. 3. Independent claim 27 states that the domain of the first speech recognizer "comprises at least a first language," the "second domain comprises at least a second language," and "the second speech recognizer is a multi-lingual speech recognizer." Id. at col. 14, ll. 20-24.

The parties agree that "language" has a plain and ordinary meaning. See Tr. at 20:3-9 (statement by Nuance's counsel at Markman hearing); Dkt. 91-1 at 6 (audio narration of Omilia's technical tutorial) ("A language is a collection of words and a grammar which allows members of a community to communicate fluently. A language also includes pronunciation rules which express words in terms of phones.").

Omilia urges the Court to adopt a definition that would prohibit the second language from being incorporated in the first speech recognizer. No such limitation exists in the claim language or in the specification, however. Claims 12, 25, and 27 all state that the first domain includes "at least a first

language." '925 patent at col. 11, ll. 59-60, col. 13, ll. 30-31; col. 14, ll. 20-21. The inclusion of the term "at least" makes clear that components of languages other than the first language may be present in the first speech recognizer. See Howmedica Osteonics Corp. v. Wright Med. Tech., Inc., 540 F.3d 1337, 1344 (Fed. Cir. 2008) (noting "at least one" means "one or more"); see also '925 patent at col. 6, ll. 3-8 (explaining that in the specification a "domain might refer to a certain language, a multitude of languages, a dialect or set of dialects, [or] a certain task area or set of task areas").

Omilia argues that the second language cannot be incorporated into the first speech recognizer because the specification states that the '925 method "can be used for the . . . incorporation of a new language into a true multi-lingual speech recognizer." '925 at col. 9, ll. 19-24 (emphases added). Omilia contends that, because this portion of the specification frames the second language as a "new language," the term "a second language" must be construed as a language not incorporated into the first domain. However, this portion of the specification merely states that the '925 method "can be" used to add a new language. A court should "not read limitations from the specification into claims" absent indication that the patentee intended the claims to be so limited. Thorner v. Sony Computer Entm't Am. LLC, 669 F.3d 1362, 1366 (Fed. Cir. 2012);

see also Liebel-Flarsheim Co. v. Medrad, Inc., 358 F.3d 898, 904 (Fed. Cir. 2004) ("[I]t is improper to read a limitation from the specification into the claims."). Omilia provides no such indication of intent.

The Court therefore construes "a second language" to mean "a language other than the first language of the first domain/speech recognizer."

II. "multi-lingual speech recognizer" (Claims 12, 25, 27)

Claim Term	Nuance's Proposed Construction	Omilia's Proposed Construction
"multi-lingual speech recognizer"	A speech recognizer with an acoustic model that covers sound units of multiple languages	A speech recognizer incorporating at least the first language and adding the second language (i.e. not present/incorporated in the first recognizer)

Nuance argues that a "multi-lingual" speech recognizer has an acoustic model that covers "sound units" of multiple languages. However, the claims refer only to "language[s]." See '925 patent at col. 11, ll. 59-64; col. 13, l. 29-col. 14, l. 3; col. 14, ll. 20-24. Omilia argues that a multi-lingual speech recognizer should "incorporat[e]" a first language and "add[]" a second language, but for the reasons stated above, claims 12, 25, and 27 do not use those terms. The Court will not import limitations from the specification into claim terms, absent a

clear indication of intent by the patentee. See Liebel-Flarsheim Co., 358 F.3d at 904.

The Court therefore gives "multi-lingual" its plain and ordinary meaning. See Dkt. 85-5 (Merriam-Webster Dictionary) ("of, having, or expressed in several languages" or "using or able to use several languages especially with equal fluency"); see also Dkt. 85-6 (Dictionary.com) ("dealing with or involving several or many languages"). "Multi-lingual speech recognizer" accordingly means "a speech recognizer that is able to cover more than one language."

III. "generating a second acoustic model" (Claim 27)

Claim Term	Nuance's Proposed Construction	Omilia's Proposed Construction
"generating a second acoustic model"	Plain and ordinary meaning	"generating a second acoustic model" requires "re-estimating said first decision network and said corresponding first phonetic contexts"

The parties dispute whether independent claim 27 requires the computerized generation of a second speech recognizer to occur via re-estimation of the first decision network and its corresponding phonetic contexts. Omilia contends that the specification language and prosecution history demonstrate that "generating a second acoustic model," as used in claim 27, requires "re-estimating said first decision network and said

corresponding first phonetic contexts." Nuance argues that the inclusion of the "re-estimating" limitation in claims 1 and 14 demonstrates that its omission in claim 27 was intentional. Moreover, Nuance argues that Omilia's proposed construction would impermissibly import limitations from the specification into claim 27.

The Court may depart from the plain and ordinary meaning of a claim term if "the patentee disavows the full scope of the claim term either in the specification or during prosecution." Golden Bridge, 758 F.3d at 1365. Disavowal "need not be explicit." Poly-America, L.P. v. API Indus., Inc., 839 F.3d 1131, 1136 (Fed. Cir. 2016). "[A]n inventor may disavow claims lacking a particular feature when the specification describes 'the present invention' as having that feature." Id. "An inventor may also disavow claim scope 'by distinguishing the claimed invention over the prior art.'" Techtronic Indus. Co. Ltd. v. Int'l Trade Comm'n, 944 F.3d 901, 907 (Fed. Cir. 2019) (quoting Ekchian v. Home Depot, Inc., 104 F.3d 1299, 1304 (Fed. Cir. 1997)).

Here, the '925 patent disavows the full scope of the claim term "generating an acoustic model." After describing prior art, the specification states, "Orthogonally to these previous approaches, the present invention focuses on the re-estimation of phonetic contexts." '925 patent at col. 6, ll. 66-67; see

Regents of Univ. of Minn. v. AGA Med. Corp., 717 F.3d 929, 936 (Fed. Cir. 2013) (explaining that when a patent "describes the features of the 'present invention' as a whole, this description limits the scope of the invention"); Hill-Rom Servs., Inc. v. Stryker Corp., 755 F.3d 1367, 1372 (Fed. Cir. 2014) (stating that a patent claim is limited "when the specification describe[s] that feature as a 'very important feature . . . in an aspect of the present invention' and disparage[s] alternatives to that feature"). Furthermore, under a subsection describing the '925 patent's proffered "Solution," the specification states that the invention "is achieved by using the speech recognizer's decision network and its corresponding phonetic contexts as a starting point and by re-estimating the decision network and phonetic contexts based on domain-specific training data." '925 patent at col. 6, ll. 16-20; see also '925 patent Abstract ("A second acoustic model . . . can be generated by re-estimating the first decision network and the corresponding first phonetic contexts . . ."); id. at col. 2, ll. 40-44 (under "Summary of the Invention") ("A second acoustic model . . . can be generated by re-estimating the first decision network and the corresponding first phonetic contexts based on domain-specific training data.")

This construction is consistent with the fact that the patent describes other methods of "generating a second acoustic

model" only to distinguish those methods from the '925 method. See '925 patent at col. 1, l. 66-col. 2, l. 3; col. 5, ll. 27-46; see also Dkt. 86 ¶ 49 (Nuance's expert declaration) (asserting that a second acoustic model could be generated without re-estimating but failing to state whether those other methods predated the '925 patent).

The patentee also disavowed the full scope of the claim term during prosecution. The original patent application included claim 27 as a dependent claim that depended upon an independent claim specifying that the generation of a second acoustic model occurred "by re-estimating said first decision network and said corresponding first phonetic contexts." See Dkt. 84-4 at 24, 22 (Nov. 13, 2001). The Examiner later rejected the underlying independent claim and directed the Applicants to rewrite claim 27 in independent form "including all of the limitations of the base claim and any intervening claims." Dkt. 84-3 at 9 (May 23, 2005). Cf. 3M Innovative Props. Co. v. Tredegar Corp., 725 F.3d 1315, 1332 (Fed. Cir. 2013) (explaining that "an examiner's statement during reexamination [is] . . . representative of how one of skill in the art would understand the term."). The Applicants then proposed an independent replacement for the previously dependent claim 27. See Dkt. 84-5 at 8 (Feb. 3, 2005). The Applicants failed to include the re-estimation language but did explain that the newly written claim

"emphasize[d] the subject matter that [was already] indicated as allowable" by the Examiner. Id. at 11. The Applicants also stated that "[n]o new matter [was] added as a result of [the] amendment[]." Id.

Nuance contends that the exclusion of the "re-estimating" limitation from claim 27 at that time was intentional because the Applicants added the "re-estimating" language to other claims during the same editing step. Cf. Akzo Nobel Coatings, Inc. v. Dow Chem. Co., 811 F.3d 1334, 1340 (Fed. Cir. 2016) (rejecting proposed construction that would render a limitation superfluous). However, Applicants represented that they had added no new matter at that step. Accordingly, while the prosecution history is not crystal clear, the patentee's statements undercut the assertion that "generating a second acoustic model" should be broadly read. Read as a whole, the prosecution history is consistent with the construction of "generating a second acoustic model" as requiring "re-estimating said first decision network and said corresponding first phonetic contexts."

Nuance also argues that Omilia's proposed construction either impermissibly gives "generating a second acoustic model" a different meaning in claim 27 than in claims 1, 2, 14, and 15, or renders the "re-estimating" limitation in those claims redundant. The Federal Circuit recognizes a "presumption that

the same terms appearing in different portions of the claims should be given the same meaning," but this presumption can be defeated if "it is clear from the specification and prosecution history that the terms have different meanings at different portions of the claims." PODS, Inc. v. Porta Stor, Inc., 484 F.3d 1359, 1366 (Fed. Cir. 2007) (quoting Fin Control Sys. Pty, Ltd. v. OAM, Inc., 265 F.3d 1311, 1318 (Fed. Cir. 2001)). Here, Omilia has produced clear language in the specification and prosecution history that "generating a second acoustic model" has a different meaning in claim 27 than in claims 1, 2, 14, and 15.

The Court holds that the term "generating a second acoustic model" in claim 27 requires "re-estimating said first decision network and said corresponding first phonetic contexts."

IV. "automatically generate/ing" (Claims 1, 2, 14, 15)

Claim Term	Nuance's Proposed Construction	Omilia's Proposed Construction
"automatically generate/ing"	Generate/ing, at least in part by a computer	Generate/ing by a computer without human intervention

Claims 1 and 2 provide "computerized method[s] of automatically generating from a first speech recognizer a second speech recognizer." '925 Patent at col. 10, ll. 42-43; col. 10, ll. 62-63. Claims 14 and 15 claim a "machine-readable storage, having stored thereon a computer program having a plurality of

code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer." Id. at col. 12, ll. 1-5; col. 12, ll. 25-29.

The parties agree that the claimed steps in claims 1, 2, 14, and 15 must be performed by a computer. However, they disagree as to whether the term "automatically generate/ing" permits or excludes human intervention as to other unclaimed steps during the execution of the claimed method.

Nuance argues that "automatically generate/ing" means "generate/ing at least in part by a computer." This construction encompasses embodiments that might require human intervention during unclaimed patent steps. In support, Nuance identifies two unclaimed portions of the patent that contemplate human activity, namely the collection of training data, '925 Patent at col. 7, ll. 32-36, col. 9, ll. 45-47, and the setting of "[p]hone questions and splitting thresholds," id. at col. 7, ll. 66-col. 8, l. 3. Neither of these manual activities is recited in the relevant claims. Nuance's expert states that because "the claimed invention does not exclude manual steps . . . a human may initiate, interrupt, or provide input for the computerized methods or apparatuses." Dkt. 86 ¶¶ 54, 59.

In contrast, Omilia argues that automatically generating means generating "by a computer without human intervention." Omilia contends that because claims 1 and 2 already provide for

a "computerized method" and claims 14 and 15 provide for "a computer program . . . executable by machine," the patent's inclusion of the phrase "automatically generate/ing" must preclude human intervention, even during unclaimed steps that occur during the execution of the claimed method. In support, Omilia points to a portion of the specification that states that the "present invention also can be realized . . . [in] [a]ny kind of computer system." '925 patent at col. 3, ll. 9-11.

The Court adopts the plain and ordinary meaning of "automatically generate/ing." The Federal Circuit has used "automatic" to mean "a process that, once initiated, is performed by a machine without the need for manually performing that process." See Whitserve, LLC v. Comput. Packages, Inc., 694 F.3d 10, 19 (Fed. Cir. 2012). This construction of "automatic" does not preclude all human intervention. See id. ("A machine may still perform the claimed process automatically, even though a human might manually initiate or interrupt the process."); see also CollegeNet, Inc. v. ApplyYourself, Inc., 418 F.3d 1225, 1235 (Fed. Cir. 2005) (construing "automatically" to mean "once initiated, the function is performed by a machine, without the need for manually performing the function").

ORDER

For the reasons stated above, the Court construes the disputed terms as follows:

- "a second language" means "a language other than the first language of the first domain/speech recognizer"
- "multi-lingual speech recognizer" means "a speech recognizer that is able to use more than one language"
- "generating a second acoustic model" requires "re-estimating said first decision network and said corresponding first phonetic contexts"
- "automatically generate/ing" means "generating by a computer without the need for manually performing that process but not excluding human intervention"

SO ORDERED.

/s/ PATTI B. SARIS
Patti B. Saris
United States District Judge